LTL-based Verification of Reconfigurable Workflows

Manuel Mazzara
Innopolis University, Russia

arXiv:1607.03649v1 [cs.SE] 13 Jul 2016

Abstract

Logics and model-checking have been successfully used in the last decades for modeling and verification of various types of hardware (and software) systems. While most languages and techniques emerged in a context of monolithic systems with a limited self-adaptability, modern systems require approaches able to cope with dynamically changing requirements and emergent behaviors. The emphasis on system reconfigurability has not been followed by an adequate research effort, and the current state of the art lacks logics and model checking paradigms that can describe and analyze complex modern systems in a comprehensive way. This paper describes a case study involving the dynamic reconfiguration of an office workflow. We state the requirements on a system implementing the workflow and its reconfiguration and we prove workflow reconfiguration termination by providing a compilation of generic workflows into LTL, using the Bound model checker Zot. The objective of this paper is demonstrating how temporal logics and model checking are effective in proving properties of dynamic, reconfigurable and adaptable systems. This simple case study is just a "proof of concept" to demonstrate the feasibility of our ideas.

Keywords: Workflow, LTL, Reconfiguration, Model checking

1 Introduction

Dependable systems have been investigated for decades [3], being software reliability of primary importance not only (and not in particular) for systems of everyday use, like operating systems or software for text and video editing or music players, but especially for life-critical systems, like those in automotive and aviation systems. However, modern systems impose new challenges to system and software engineers as well as to the overall computer science community. These days we increasingly build systems that respond to external stimuli and change their computational behavior accordingly. In such

situations, their reconfigurations are externally determined. This poses reconfiguration and adaptability requirements and the need to evaluate system dependability from a broader perspective than before.

The next generation of dependable systems is expected to have significant evolution requirements [8] and it is highly likely that the system will have to be redesigned (i.e. reconfigured) during its lifetime (possibly more than once) in order to meet new requirements. Therefore, it will be impossible to foresee all the requirements that a system will have to meet in future when it is being designed [17]. In particular, certain classes of dependable systems, such as control systems, must be dynamically reconfigured [11], because it is unsafe or impractical or too expensive to do otherwise. The dynamic reconfiguration of a system is defined as the change at runtime of the structure of the system – consisting of its components and their communication links – or the hardware location of its software components [7].

This paper focuses on dynamic software reconfiguration, because software is much more mutable than hardware [16, 2]. However, several techniques that have been applied to hardware verification can still be applied to software. Logics and model-checking, for example, have been successfully used in the last decades for modeling and verification of various types of hardware (and software) systems. While most languages and techniques emerged in a context of monolithic systems with a limited self-adaptability, modern systems require approaches able to cope with dynamically changing requirements and emergent behaviors. The emphasis on system reconfigurability has not been followed by an adequate research effort, and the current state of the art lacks logics and model checking paradigms that can describe and analyze complex modern systems in a comprehensive way.

In this work, we tackle the open problem of verifying reconfigurable systems (and in particular reconfigurable workflows starting from [13]) by temporal logics and model checking. This paper describes a case study involving the dynamic reconfiguration of an office workflow similar to the one studied in [15]. The role of temporal logics in verification and validation is two-fold. First, temporal logic allows abstract, concise and convenient expression of required properties of a system. Linear Temporal Logic (LTL) is often used with this goal in the verification of finite-state models, e.g., in model checking [4]. Second, temporal logic allows a descriptive approach to specification and modeling (see, e.g., [18, 10]). A descriptive model is based on axioms, written in some (temporal) logic, defining a system by means of its general properties, rather than by an operational model based on some kind of ma-

chine behaving in the desired way. In this case, verification typically consists of satisfiability checking of the conjunction of the model and of the (negation of) its desired properties. An example of the latter approach is Bounded Satisfiability Checking (BSC) [21], where Metric Temporal Logic (MTL) specifications on *discrete* time and properties are translated into Boolean logic, in an approach similar to Bounded Model Checking of LTL properties of finite-state machines.

The rest of the paper is organized as follows: Section 2 describes the case study, in which a simple office workflow for order processing is reconfigured. The reconfiguration of the workflow is modelled and analyzed using temporal logic and, in particular, in Section 3 a general encoding of workflows into LTL is provided. In Section 4 the implementation of this translation is illustrated and tests have been carried out to validate its correctness. Finally, Section 5 draw conclusive remarks and focus on future developments of this work.

2 Case study

The case study approached in this paper describes dynamic reconfiguration of an office workflow for order processing that is commonly found in large and medium-sized organizations [9]. These workflows typically handle large numbers of orders. Furthermore, the organizational environment of a workflow, can change in structure, procedures, policies and legal obligations in a manner unforeseen by the original designers of the workflow. Therefore, it is necessary to support the unplanned change of these workflows. Furthermore, the state of an order in the old configuration may not correspond to any state of the order in the new configuration. These factors, taken in combination, imply that instantaneous reconfiguration of a workflow is not always possible; neither is it practical to delay or abort large numbers of orders because the workflow is being reconfigured. The only other possibility is to allow overlapping modes for the workflow during its reconfiguration.

The workflow consists of a graph of activities and the configuration is the structure of this graph. Initially, the workflow executes in Configuration 1 and must satisfy the requirements on Configuration 1. Subsequently, the workflow must be reconfigured through a process to Configuration 2 in a way that the requirements on the reconfiguration process and on Configuration 2 are satisfied. The full details of the two configurations and the relative requirements are given in [15]. The initial configuration of the workflow is

depicted in Figure 1.

After some time, the management of the organization using the workflow decides to change it in order to increase opportunities for sales, improve the synchronisation between `Billing` and `Shipping`, and to simplify the workflow. The new configuration of the workflow is Configuration 2 as depicted in Figure 2. The two pictures below represent a reconfiguration between the two workflows. Dashed rectangles are the portions of the workflows that are translated in Section 3.

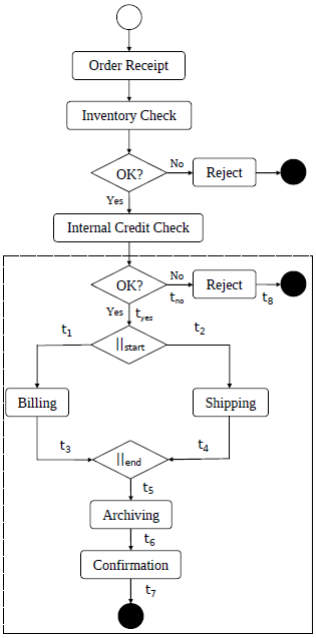

Figure 1: First configuration

In order to achieve a smooth transition from Configuration 1 to Configuration 2 of the workflow, the process of reconfiguration must meet the following requirements:

1. Reconfiguration of a workflow should not necessarily result in the rejection of an order

2. Any order being processed that was accepted **before** the start of the reconfiguration must satisfy all the requirements on Configuration 1

3. Any order accepted **after** the start of the reconfiguration must satisfy all the requirements on Configuration 2

4. The reconfiguration process must terminate

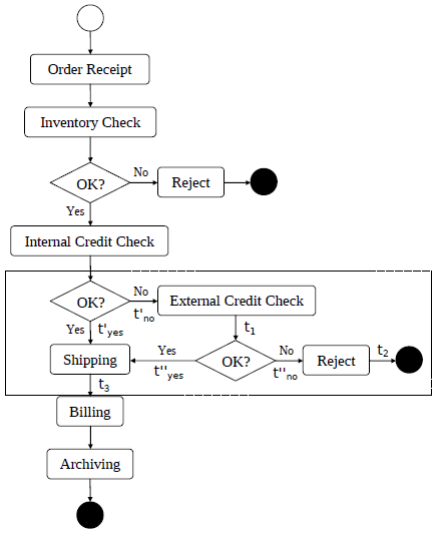

Figure 2: Second configuration

Methodology

The process of verification is depicted in Figure 3. In this work, we provide an instrument to designers for workflows revision, i.e. a procedure to follow until the requirements are finally met. To do this, we encode the system by means of a formal language, as discussed in Section 3, and we automatically determine its correctness as shown in Section 4. Differently from several others formal approaches that only offer languages without a method to apply them (for a detailed discussion see [14]), this methodology, together with software tools, aims at offering a complete practical toolkit for software and systems engineers working in the field of workflow reconfiguration.

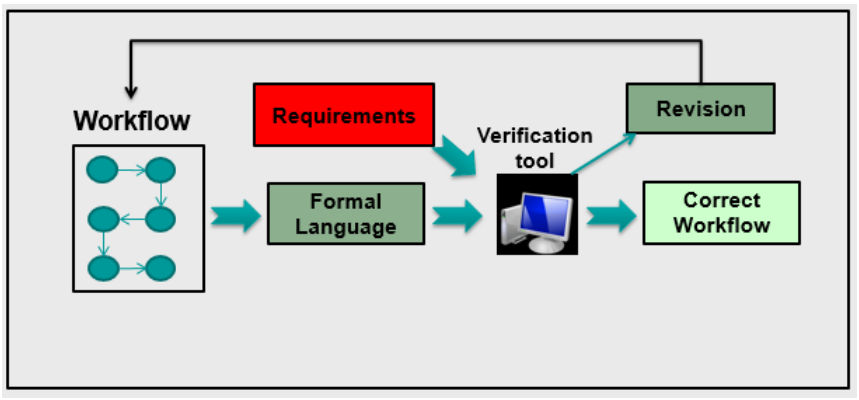

Figure 3: Workflow Verification Process

3 Verification by LTL specification

LTL [12] is one of the most popular descriptive languages for defining temporal behaviors that are represented as sequences of observations. The time model adopted in this logic is a totally ordered set (e.g., $(\mathbb{N}, <)$) whose elements are the positions where the behavior is observed. LTL allows the expression of positional orders of events both towards the past and the future. Let AP be a finite set of atomic propositions. Well-formed LTL formulae are defined as follows:

$$\phi := p \mid \phi \wedge \phi \mid \neg\phi \mid \mathbf{X}(\phi) \mid \mathbf{Y}(\phi) \mid \phi \mathbf{U} \phi \mid \phi \mathbf{S} \phi$$

where $p \in AP$, $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{U}$ and $\mathbf{S}$ are the usual "next", "previous", "until" and "since" modalities. The dual operator "release" $\mathbf{R}$ is defined as usual, i.e., $\phi \mathbf{R} \psi$ is $\neg(\neg\phi \mathbf{U} \neg\psi)$ and the "Globally" $\mathbf{G}\phi$ operator is $\mathit{false}\mathbf{R}\phi$.

The semantics of LTL formulae is defined with respect to a strict linear order representing time $(\mathbb{N}, <)$. Truth values of propositions in AP are defined by function $\pi : \mathbb{N} \to \wp(AP)$ associating a subset of the set of propositions with each element of $\mathbb{N}$. The semantics of a LTL formula ϕ at instant $i \geq 0$ over a linear structure π is recursively defined as in Table 1, A formula $\phi \in$ CLTL is *satisfiable* if there exists a pair π such that $\pi, 0 \models \phi$.

Workflows model execution of systems as sequences of activities. Transitions, conditional choices and parallel "split" interleave the activities and determine uniquely the flow of the execution, i.e., the sequence of activities that realizes the computation. An activity is an abstraction of a compound

$$
\begin{aligned}
\pi, i \models p &\Leftrightarrow p \in \pi(i) \text{ for } p \in AP \\
\pi, i \models \mathbf{X}(\phi) &\Leftrightarrow \pi, i+1 \models \phi \\
\pi, i \models \mathbf{Y}(\phi) &\Leftrightarrow \pi, i-1 \models \phi \wedge i > 0 \\
\pi, i \models \phi \mathbf{U} \psi &\Leftrightarrow \exists\, j \geq i : \pi, j \models \psi \ \wedge \pi, n \models \phi \ \forall\ i \leq n < j \\
\pi, i \models \phi \mathbf{S} \psi &\Leftrightarrow \exists\, 0 \leq j \leq i : \pi, j \models \psi \ \wedge \pi, n \models \phi \ \forall\ j < n \leq i
\end{aligned}
$$

Table 1: Semantics of CLTL. Boolean connectives are omitted for brevity.

of actions that are performed by the real systems. Although they can be modelled as atomic computations, we adopt a different perspective for which the activities, being actions in the real world, have a non-punctual duration. To translate workflows into an LTL formula, we assume that an activity is always followed by a transition, and viceversa, and that conditional tests and splits are special activities that have punctual duration. When an activity is performed, the firing of the outgoing transition lets the system change allowing it to execute the next activity. Therefore, our LTL translation allows modelling of workflows as sequences of activities and transitions, in strict alternation.

With no loss of generality, we assume that no element in the same graph is duplicated. By this assumption, we associate each element with an atomic proposition that uniquely identifies it. Observe, however, that there could be some activity that are shared by the two configurations but each transition has a unique name and no transition with the same label appears twice. We write A to represent the activity A in the workflow, while we write t to indicate a transition between two activities. If A holds at position i then the workflow is performing activity A at that position; similarly for t. We introduce $\|$ and $\oplus$ to indicate the split and the conditional activity, respectively; *start* and *end* to indicate two special activities representing the starting and the final activity of the workflow. Workflow diagrams are translated according to rules in Table 2.

Let A be an activity, $out(A)$ be the set of outgoing transition starting from A and $in(A)$ be the set of ingoing transition leading to A. By definition, a correct workflow is such that $out(A) \geq 1$, for all activity A, except for activity *end* and that $out(A) \geq 1$, for all activity A, except for activity *start*. Let $t_{\text{out}}(A)$ be the disjunction $\bigvee_{t_i \in out(A)} t_i$ and $t_{\text{in}}(A)$ be the disjunction $\bigvee_{t_i \in in(A)} t_i$.

$[A]_{t_{out}(A)}$	$A \Rightarrow (A \wedge \neg t_{\text{out}}(A))\mathbf{U}(t_{\text{out}}(A))$ (1) $\bigwedge_{i \in out(A)} (t_i \Rightarrow \mathbf{Y}(A) \wedge \neg A)$ (2)	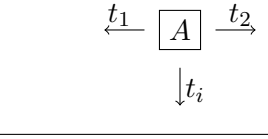
${}_{t_{in}(A)}[A]$	$A \Rightarrow (A \wedge \neg t_{\text{in}}(A))\mathbf{S}(t_{\text{in}}(A))$ (3) $\bigwedge_{i \in in(A)} (t_i \Rightarrow \mathbf{X}(A) \wedge \neg A)$ (4)	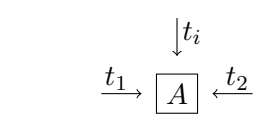
$[\cdot + \cdot]$	$t_1 \Rightarrow \neg t_2$ (5) $\oplus \Rightarrow \neg\mathbf{Y}(\oplus) \wedge \neg\mathbf{X}(\oplus)$ (6)	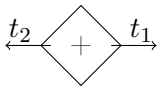
$[\cdot \mid \cdot]$	$\bigwedge_{t_1,t_2 \in out(\|)} (t_1 \Leftrightarrow t_2)$ or $\bigwedge_{t_1,t_2 \in in(\|)} (t_1 \Leftrightarrow t_2)$ (7) $\| \Rightarrow \neg\mathbf{Y}(\|) \wedge \neg\mathbf{X}(\|)$ (8)	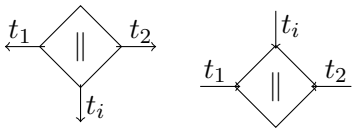

Table 2: Workflow LTL encoding

Formula (1) states that if activity A holds at the current position, then it is true, at that position, that A but not $t_{\text{out}}(A)$ holds until at least one transition in $out(A)$ holds. In this way, activity A lasts until one of its outgoing transition fires. Formula (2) imposes that if transition $t_i \in out(A)$ holds at position i, then in the previous one $i-1$, activity A holds. This necessary condition enforces that a transition fires only if the activity from which it originates has just been terminated.

Formula (3) states that if activity A holds at the current position, then it is true, at that position, that A but not $t_{\text{in}}(A)$ holds since at least one transition in $in(A)$ has been fired. In this way, activity A has lasted since one of its ingoing transition fired. Formula (4) imposes that if transition $t_i \in in(A)$ holds at position i, then, in the next one $i+1$, activity A holds. This necessary condition enforces that a transition fires only if the activity to which it leads will be performed in the next position.

The translation for $[A]_{t_{out}(A)}$ and ${}_{t_{in}(A)}[A]$ does not impose constraints on transitions in $t_{out}(A)$ and $t_{in}(A)$. In the case of $[A]_{t_{out}(A)}$, the execution flow after an activity A is nondeterministically defined as it is determined by at least one transition that fires when A terminates. Simmetrically, the

execution of an activity A that is reached by more than one transition in ${}_{t_{in}(A)}[A]$ can be started by at least one execution that is active before A. The conditional block $[\cdot + \cdot]$ and the split $[\cdot \mid \cdot]$ are modelled by using the rules (1)-(4) and, in addition, specific constraints to enforce the proper flow of execution. To model the flow of the conditional block, we force the execution of the two branches to be exclusive as for the *if-then-else* construct. The translation for the split block is similar yet it enforces the synchronization of all the transitions starting from, or yielding to, the activity split.

The conditional block $[\cdot \mid \cdot]$ is translated compositionally. For each conditional block we introduce a new fresh atomic proposition $\oplus$. Formulae (1), (2), (3) and (4) rule the semantics of the block as sequence of actions. In addition, Formula (5) impose that only one branch is executed, by forcing the strict complementarity between t_1 and t_2. Formula (6) enforces punctuality of $\oplus$ and states that if $\oplus$ holds at position i then in the next and in the previous positions it does not hold.

The translation for the split block is similar to the one defined for the conditional block. For each split block we introduce a new fresh atomic proposition $\|$. Formulae (1), (2), (3) and (4) rule the semantics of the block as sequence of actions and Formula (6) enforces punctuality, similarly to the previous case. The only difference is Formula (7) that is divided into two parts that are used exclusively. Both of them impose strict contemporaneity of all the transitions involved in activity $\|$. The first one is defined only for the initial point of a split block and states that all the outgoing transitions starting from it, occur at the same time. The second one is similar but for the final point of a split block, where all the parallel computations must join before proceeding further. It states that all the ingoing transitions leading to it, occur at the same time.

For the sake of brevity, we provide, as an example, the translation of a portion of the two workflows of Figure 1 and 2. Let us consider the first one 1 from the second conditional. We introduce the following atomic propositions representing the activities. We use $\oplus$ for modelling the conditional block, *Rej* for *Reject*, $\|_{start}$ and $\|_{end}$ for the activities defining the starting and the ending point of the join block, *Arch*, *Ship*, *Bill* and *Conf* for *Archiving*, *Shipping*, *Billing* and *Confirmation*. Finally, we introduce the final state *end*. The following formulae model the portion of the workflow. We unify the translation for activities *Ship*, *Ship*, *Bill* and *Conf*, being similar. Let $A \in \{Ship, Arch, Bill, Conf, Rej\}$ and $t' \in \{t_1, t_2, t_5, t_6, t_{no}\}$, $t'' \in \{t_3, t_4, t_6, t_7, t_8\}$. We use the following cases for (A, t', t''): $(Ship, t_2, t_4)$,

$(Bill, t_1, t_3)$, $(Arch, t_5, t_6)$, $(Conf, t_6, t_7)$ and (Rej, t_{no}, t_8).

$$\oplus \Rightarrow (\oplus \wedge \neg(t_{yes} \vee t_{no}))\mathbf{U}(t_{yes} \vee t_{no})$$
$$t_{yes} \Rightarrow \mathbf{Y}\,(\oplus) \wedge \neg\oplus$$
$$t_{no} \Rightarrow \mathbf{Y}\,(\oplus) \wedge \neg\oplus$$
$$t_{yes} \Rightarrow \neg t_{no}$$
$$\oplus \Rightarrow \neg\mathbf{Y}\,(\oplus) \wedge \neg\mathbf{X}\,(\oplus)$$

$$\|_{start} \Rightarrow (\|_{start} \wedge \neg t_{yes})\mathbf{S}\; t_{yes}$$
$$t_{yes} \Rightarrow \mathbf{X}\,(\|_{start}) \wedge \neg\|_{start}$$
$$\|_{start} \Rightarrow (\|_{start} \wedge \neg(t_1 \vee t_2))\mathbf{U}(t_2 \vee t_2))$$
$$t_2 \Rightarrow \mathbf{Y}\,(\|_{start}) \wedge \neg\|_{start}$$
$$t_1 \Rightarrow \mathbf{Y}\,(\|_{start}) \wedge \neg\|_{start}$$
$$t_1 \Leftrightarrow t_2$$
$$\|_{start} \Rightarrow \neg\mathbf{Y}\,(\|_{start}) \wedge \neg\mathbf{X}\,(\|_{start})$$

$$A \Rightarrow (A \wedge \neg t')\mathbf{S}\; t'$$
$$t' \Rightarrow \mathbf{X}\,(A) \wedge \neg A$$
$$A \Rightarrow (A \wedge \neg t'')\mathbf{U}\; t'')$$
$$t'' \Rightarrow \mathbf{Y}\,(A) \wedge \neg A$$

$$\|_{end} \Rightarrow (\|_{end} \wedge \neg(t_3 \vee t_4))\mathbf{S}\; (t_3 \vee t_4)$$
$$t_4 \Rightarrow \mathbf{X}\,(\|_{end}) \wedge \neg\|_{end}$$
$$t_3 \Rightarrow \mathbf{X}\,(\|_{end}) \wedge \neg\|_{end}$$
$$\|_{end} \Rightarrow (\|_{end} \wedge \neg t_5)\mathbf{U}\; t_5)$$
$$t_5 \Rightarrow \mathbf{Y}\,(\|_{end}) \wedge \neg\|_{end}$$
$$t_3 \Leftrightarrow t_4$$
$$\|_{end} \Rightarrow \neg\mathbf{Y}\,(\|_{end}) \wedge \neg\mathbf{X}\,(\|_{end})$$

$$end \Rightarrow (end \wedge \neg(t_7 \vee t_8))\mathbf{S}\; (t_7 \vee t_8)$$
$$t_7 \Rightarrow \mathbf{X}\,(end) \wedge \neg end$$
$$t_8 \Rightarrow \mathbf{X}\,(end) \wedge \neg end$$

The partial translation of the second workflow follows. In particular, only the second and the third conditional block and the activities *External Credit Check* (*Ecc*) and *Shipping* are considered. We represent the first conditional with $\oplus'$ and the second one with $\oplus''$. We abbreviate the notation for the two blocks with $\oplus^*$, where $*$ is the a quote $'$ or a double quote $''$, and for the associated translations with t^*_{yes} and t^*_{no}. The translation for $\oplus^i$ follows exactly the same structure as the one previously defined for the first workflow, so we do not provide it.

$$Ecc \Rightarrow (Ecc \wedge \neg t'_{no})\mathbf{S}\; t'_{no}$$
$$t'_{no} \Rightarrow \mathbf{X}\,(Ecc) \wedge \neg Ecc$$
$$Ecc \Rightarrow (Ecc \wedge \neg t_1)\mathbf{U}\; t_1)$$
$$t_1 \Rightarrow \mathbf{Y}\,(Ecc) \wedge \neg Ecc$$

$$Ship \Rightarrow (Ship \wedge \neg(t'_{yes} \vee t''_{yes}))\mathbf{S}\; (t'_{yes} \vee t''_{yes})$$
$$t'_{yes} \Rightarrow \mathbf{X}\,(Ship) \wedge \neg Ship$$
$$t''_{yes} \Rightarrow \mathbf{X}\,(Ship) \wedge \neg Ship$$
$$Ship \Rightarrow (A \wedge \neg t_3)\mathbf{U}\; t_3)$$
$$t_3 \Rightarrow \mathbf{Y}\,(t_3) \wedge \neg Ship$$

$$Rej \Rightarrow (Rej \wedge \neg t''_{no})\mathbf{S}\; t''_{no}$$
$$t''_{no} \Rightarrow \mathbf{X}\,(Rej) \wedge \neg Rej$$
$$Rej \Rightarrow (Rej \wedge \neg t_2)\mathbf{U}\; t_2)$$
$$t_2 \Rightarrow \mathbf{Y}\,(Rej) \wedge \neg Rej$$

In this work, we assume that when a workflow instance terminates, it never resumes. We enforce it through the formula $end \Rightarrow \mathbf{X}\,(\mathbf{G}\,(\neg end))$.

The previous formulae only apply for the case of one workflow. In our model, many workflows can be active at the same time and to represent different instances of the same workflow, we introduce as many activities A^j and transitions t^j as the number of workflows to be represented, where $j \in \mathbb{N}$ is the index representing the j-th instance. All the formulae defining the workflow are replicated to define the behaviour for all the instances of the

workflow. To model the reconfiguration process, we need to represent when an instance is active or not. A workflow is active at a certain position i when at least one of its activity or transition holds at i. We introduce for all the workflows an atomic proposition W^j, with $j \in \mathbb{N}$, that holds if, and only if, the previous condition is met. Let E_i^j be the set of propositions representing activities and transitions defining the j-th instance of the workflow for the i-th configuration. The following formula constraints W^j:

$$W^j \Leftrightarrow \bigvee_{e \in E_1^j \cup E_2^j} e. \tag{9}$$

Observe that $E_1^j \cap E_2^j$ may, in general, be non empty. This is the case of configurations of the same workflow that share some activity or transitions. A configuration of a workflow is represented as the conjunction of all the formulae that translate it, that are defined by rules in Table 3. We introduce proposition C_1^j and C_2^j to capture the configuration in Figure 1 and in Figure 2, for all the instances j. Let WS_i^j be the set of all formulae translating configuration $i \in \{0, 1\}$ of (an instance of) workflow j that are defined according to rules in Table 3. Then, the formula defining the rules characterizing configuration i is

$$C_i^j \Rightarrow \bigwedge_{\phi_h \in WS_i^j} \phi_h. \tag{10}$$

Formula10 states that if C_i^j holds at the current position, then all the formulae ϕ_h in WS_i^j determine the current configuration and rule the execution flow. Given an instance j, to guarantee that all, and only, the activities and transitions that appear in the first configuration may occur when C_1^j holds (but not the activities and transitions that appear in the second configuration yet not in the first one), we add the following formula. Let $\bar{E}_1^j = E_2^j \backslash (E_1^j \cap E_2^j)$ be the set of activities and transitions that appear in the second configuration but not in the first one; and let $\bar{E}_2^j = E_1^j \setminus (E_1^j \cap E_2^j)$ be the set of activities and transitions that appear in the first configuration but not in the first second.

$$\bigwedge_{i \in \{1,2\}} C_i^j \Rightarrow (\bigvee_{e \in E_i^j} e) \wedge \neg(\bigwedge_{e \in \bar{E}_i^j} e) \tag{11}$$

The reconfiguration is represented by proposition R. We assume that there is only one reconfiguration in the system and that, from the position

where it occurs, all the workflows behave as the second configuration defined in Figure 2. Before the reconfiguration point, instance behaviour conforms to the first workflow of Figure 1. To model the reconfiguration, being it unique, we first impose that when R holds, at a certain position, then it holds forever in the future. This is achieved by the formula $R \Rightarrow \mathbf{G}(R)$. The choice of the workflow that rules the behaviour of an active instance is determined in the following way and it is defined by the next formula. If an instance starts the execution before the point where the reconfiguration occurs then its behaviour conforms to the workflow of Figure 1. In the other case, i.e., the workflow instance starts after the reconfiguration point, then its behaviour conforms to the workflow of Figure 2.

$$\begin{array}{c} \left(\begin{array}{c} (W^j \wedge \neg R)\mathbf{U}(\neg W^j) \ \vee \\ (W^j \wedge \neg R)\mathbf{U}((W^j \wedge R)\mathbf{U}(\neg W^j \wedge R)) \end{array} \right) \Rightarrow C_1^j \quad \wedge \\ (W^j \wedge R)\mathbf{U}(\neg W^j) \Rightarrow C_2^j \quad \wedge \\ \neg \mathbf{F}\left(W^j \wedge \neg R\right) \Rightarrow \mathbf{G}(C_2^j). \end{array} \tag{12}$$

The first formula of the previous conjunction states that if, at some position, (i) the instance j is active when the reconfiguration has not yet occurred until it terminates, (ii) the instance j is active when the reconfiguration has not yet occurred until a position where j is still active and the reconfiguration occurred until it terminates, then its behaviour is defined by the first configuration of the workflow in Figure 1. The second formula, defines the opposite situation where an instance j starts after the reconfiguration, i.e., the instance j is active when the reconfiguration has occurred until it terminates. In such case, the second workflow of Figure 2 defines the behaviour for j. The last formula imposes the behaviour defined by the second workflow over the whole time line if the instance j never start before the reconfiguration. The definition of the model is the conjunction of all the previous formulae, i.e., the translations of the workflows, Formula (9), Formula (10) and Formula (12), globally quantified over time by $\mathbf{G}$.

4 Experimental Results

In this section, we briefly illustrate the implementation of the translation defined in Section 3.The test has been carried out to validate the correctness of rules in Table 2 from a practical point of view and it is, at the current

state, only a proof-of-concept. We considered two workers, realizing the workflow of Figure 1 and 2, to prove the reachability of final states when a reconfiguration occur. In particular, we impose that the reconfiguration process is activated 3 positions of time after the origin and that the first worker starts the execution at position 2 while the second one at position 4.

Checking the satisfiability of the translation of the workflow allowed us to verify whether there exists one execution with reconfiguration that satisfies the desired behaviour. In other words, this means that the process of reconfiguration is correctly performed by the model. To check the satisfiability of our LTL model and the property of termination, we exploit the Bounded Satisfiability Checking (BSC) [21] approach. The key idea behind the BSC is to build a finite representation, of length k, of an infinite ultimately periodic LTL model of the form $\alpha\beta^{\omega}$, where α and β are finite words over the alphabet 2^{AP}. Although the standard technique to prove the satisfiability of LTL formulae is based on the construction of Büchi automata [24], the evidence has turned out that it is rather expensive in practice, even in the case of LTL (the size of the automaton is exponential with respect to the size of the formula) and that new techniques should be investigated. BSC tackles the complexity of checking the satisfiability for LTL formulae by avoiding the unfeasible construction of the whole automaton. By unrolling the semantics of the formula for a finite number $k > 0$ of steps, BSC tries to build a bounded representation $\alpha\beta$, of length k, of an infinite ultimately periodic model for the formula of the form $\alpha\beta^{\omega}$. [5] proves that BSC problem for LTL and its extension is complete and that it can be reduced to a decidable Satisfiability Modulo Theory (SMT) problem. It turns out that, in many cases, as our tests confirm, that the expensive construction of the Büchi automaton for the LTL formula can be avoided as it is possible to find finite models of small length in a very short time.

All our tests ware carried out by using the ℤot [20] toolkit. ℤot is a Bounded Model/Satisfiability checker that takes as input specifications written in a variety of temporal logics, and determines whether they are satisfiable or not. Through this basic mechanism, it can perform verification of user-defined properties for the desired models, the workflow in our case.

The methodology that we exploit to verify the translation and the reachability property for the workflow is the following. Let S be formula translating the workflow of Figure 1 or Figure 2. If S is fed to the ℤot tool "as is", ℤot will look for one of its execution; if it does not find one (i.e., if the model is unsatisfiable), then the our translation is contradictory, hence it contains

some flaws. Now, let P be a formula which formalizes the property of reachability we want to check on the workflow. If $S \wedge \neg P$ is unsatisfiable, this means that there is no execution that satisfies the workflow (i.e., S), that also satisfies $\neg P$, that is, that violates the property P. If no execution violates property P, then the latter actually holds for the workflow. If, on the other hand, $S \wedge \neg P$ is satisfiable, this means that there is at least one execution that satisfies both S and $\neg P$; that is, there is at least one execution of the workflow that violates the property, so the property does not hold. If ℤot determines that a formula is satisfiable, then the tool produces an execution that satisfies it that one can use to check the correcteness of the modelling.

ℤot performs the checks outlined previously by encoding temporal logic formulae into the input language of various solvers. In particular, ℤot supports two kinds of solvers: SAT solvers and SMT solvers. SAT solvers are capable of taking, as input, formulae written in propositional logic and determine whether they are satisfiable or not. SMT solvers do the same, but they accept, as input, formulae written in logics (fragments of First-Order Logic) that are richer than the simple propositional logic. Over the last few years, both SAT solvers and SMT solvers have made great strides in terms of their performances, so that they have become viable engines for fully automated logic-based verification approaches such as the one realized by ℤot. In addition, most SAT/SMT solvers accept inputs written in a standard format, which makes them easily interchangeable. This is very useful, since different solvers implement different heuristics, and the "best" solver does not exist in absolute terms, but only on a model-by-model case.

ℤot plugins implement the primitives that are used in ℤot scripts to define the models to be analysed. Since the plugins are implemented as Lisp modules, the primitives are essentially Lisp declarations. Then, ℤot scripts, which contain both the model to be analysed and the necessary commands to invoke the desired solver, are a collection of Lisp statements.

Table 3 shows the time required by the tool to check the reachability properties ($\mathbf{F}(s)$, where s is activity *Confirmation* or *Archiving*), the memory occupation and the result, i.e. whether the property is satisfied or not. All tests have been carried out on a 3.3 Ghz quad core PC with 16 Gbytes of Ram.

The first formula refers to the reachability of the Accept state of the workflow in Figure 1, represented by the first worker, which starts before the instant of reconfiguration. To show the impact of the model checker, we have changed the bound k, which is a user-defined parameter that corresponds to

the maximal length of runs analysed by Zot. It corresponds to the number of steps needed to build the bounded representation of the model. The values chosen are $k = 20$ and $k = 100$.

The third formula refers to the reachability of the Accept state of the workflow in Figure 2, represented by the second worker, which starts after the instant of reconfiguration. It is again been verified using two different values for the bound k.

Table 3: Test results.

Formula	Time (sec)	Memory (Mb)	Bound
Accept state, Configuration 1	24.417	164	20
Accept state, Configuration 1	82.734	309	100
Accept state, Configuration 2	28.873	187	20
Accept state, Configuration 2	94.74387	367	100

As it is expected, the time and memory consumption increases with the length of the bound of the analysed traces. Although the model is composed of a number of atomic propositions which is equal to the sum of the number of states and transitions of the two workflows (about 50), the needed time to perform the analysis is very low: this is due to the type of property analysed. For reachability properties the analysis is very faster, since the tool stops as soon as it finds a run that satisfies the formula; for example, deadlock detection analysis could take a long time, minutes or hours, as the tool must exhaustively analyse all possible runs.

Furthermore, our tests have formally demonstrated the termination of workflows, before and after reconfiguration.

We can conclude that it is feasible, using modern model checking tools such as Zot, to perform the formal verification of some simple properties, like reachability properties, in a limited amount of time. We left as future work the implementation and testing of more interesting properties, such as real-time properties.

5 Conclusions and future work

The objective of this paper is demonstrating how temporal logics and model checking are effective in proving properties of dynamic, reconfigurable and

adaptable systems, in this specific case business workflows. We proved workflow reconfiguration termination by providing a compilation of generic workflows into LTL and using ℤot as a model checker. The idea of this research stem by the observation that reconfigurability has not been thoroughly investigated in the context of verification and the current state of the art fundamentally lacks logics and model checking paradigms able to describe and analyze complex modern systems. Therefore, with this research strand we are aiming towards a full exploitation of temporal logics and model checking techniques for verification of self-adaptability, dynamically changing requirements and emergent behaviors.

This simple case study is just a "proof of concept" to demonstrate the feasibility our ideas, which now need to be challenged investigating a real industrial case study. In particular, the workflow patterns here analyzed are limited with respect to a real scenario and therefore the nature of the encoded workflow and the details of the encoding need to be fully worked out. Workflow patterns as presented in [23] need to be investigated and eventually encoded. Once workflows are intended as graphs and transitions are treated like in our encoding, similarities emerge with the Petri Nets approach. In particular, we need to compare our work with Workflow Nets [1] and exploit possible synergies.

Future work aims at extending the current translation of workflows by using more expressive logics. In particular, we would extend the basic definition of workflow to *timed* workflow by adding timing constraints on activities and transitions. This extension allows us to constraint the duration of activities and of the execution flow with respect to time bounds that may vary with respect to certain environment parameters that are measured along the execution of the workflow. To model timed workflow we will exploit CLTLoc [6], that is an LTL based logic where atomic formulae are both atomic propositions and constraints over dense clocks.

Another field of future investigation, which is in the same direction of timed workflow, is the modelization of workflows where some activities take a neglected duration time respect to other ones. Infact, there are different real-time systems where the various activities of the workflow modeling them have a huge different time duration; for example, an activity which consists of the automatic reception of orders by e-mail has – taken singularly – a very short duration respect to the activity of checking the content of those orders, provided by humans in hours or days. Due to the different order of magnitudes in the duration of the two different activities, usually the first

activity, with a neglected duration, is modeled as having a logical zero time duration. The zero-time modelization deals to a lot of different problematics, such as the rising of Zeno behaviours – a sequence of infinite actions or activities executed in a finite or zero duration time, which is not a feasible behaviour of a real system – or the fact that the "state" of the system – which is represented by the set of atomic propositions which is true in a certain instant of time – is not anymore a function of time, which is in a certain sense counterintuitive and may lead to logical contradiction. Some temporal logics adopt a time structure where more than one system state may be associated with a single time instant – an example, reported in [19], is the super-dense time model – and therefore must provide distinct notations to refer to time and state change. In these notations, the progress of time and state evolution are fully decoupled, which is rather unnatural in most practical cases. To overcome this limitations, the work in [10] introduce a new metric temporal logic called *X-TRIO*, which exploits the concepts of *Non-Standard Analysis* [22], substituting each zero-time activity with an "infinitesimal" duration time activity, enriching the underline time model. As future work, we will study the way to "glue" together the CLTLoc logic with X-TRIO, adding to the first temporal logic the capability to model neglected duration time activities provided by the latter.

Acknowledgements. I would like to acknowledge the help given by numerous colleagues, in particular: Anirban Bhattacharyya, Marcello Maria Bersani and Luca Ferrucci. This work has been supported by the Russian Ministry of education and science with the project "Development of new generation of cloudy technologies of storage and data control with the integrated security system and the guaranteed level of access and fault tolerance" (agreement: 14.612.21.0001, ID: RFMEFI61214X0001).

References

[1] Wil M. P. van der Aalst. Verification of workflow nets. In *Proceedings of the 18th International Conference on Application and Theory of Petri Nets*, ICATPN '97, pages 407–426, London, UK, UK, 1997. Springer-Verlag.

[2] Faisal Abouzaid, John Mullins, Manuel Mazzara, and Nicola Dragoni.

Modelling and analysis of dynamic reconfiguration in bp-calculus. In *KES-AMSTA*, pages 534–543, 2012.

[3] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Secur. Comput.*, 1(1):11–33, January 2004.

[4] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.

[5] Marcello M. Bersani, Achille Frigeri, Matteo Rossi, and Pierluigi San Pietro. Completeness of the bounded satisfiability problem for constraint LTL. In *Reachability Problems*, volume 6945 of *LNCS*, pages 58–71. 2011.

[6] Marcello M. Bersani, Matteo Rossi, and Pierluigi San Pietro. A tool for deciding the satisfiability of continuous-time metric temporal logic. In *Proceedings of the International Symposium on Temporal Representation and Reasoning (TIME)*, pages 99–106, 2013.

[7] A. Bhattacharyya. *Formal Modelling and Analysis of Dynamic Reconfiguration of Dependable Systems*. PhD thesis, Newcastle University School of Computing Science, 2013.

[8] L. Coyle, M. Hinchey, B. Nuseibeh, and J. L. Fiadeiro. Guest editors' introduction: Evolving critical systems. *IEEE Computer*, 43(5):28–33, 2010.

[9] Clarence Ellis, Karim Keddara, and Grzegorz Rozenberg. Dynamic change within workflow systems. In *Proceedings of Conference on Organizational Computing Systems*, COCS '95, pages 10–21, New York, NY, USA, 1995. ACM.

[10] L. Ferrucci, D. Mandrioli, A. Morzenti, and M. Rossi. A metric temporal logic for dealing with zero-time transitions. In *Proc. of 19th International Symposium on Temporal Representation and Reasoning*, pages 81–88. IEEE Computer Society, 2012.

[11] G. Karsai, F. Massacci, L. J. Osterweil, and I. Schieferdecker. Evolving embedded systems. *IEEE Computer*, 43(5):34–40, 2010.

[12] Orna Lichtenstein, Amir Pnueli, and Lenore Zuck. The glory of the past. In *Proc. of Logics of Programs*, volume 193 of *LNCS*, pages 196–218. Springer, 1985.

[13] Manuel Mazzara. *Towards Abstractions for Web Services Composition.* PhD thesis, Dipartimento di Scienze dell'Iinformazione, University of Bologna, Italy, 2006.

[14] Manuel Mazzara. Deriving specifications of dependable systems: toward a method. In *12th European Workshop on Dependable Computing (EWDC)*, 2009.

[15] Manuel Mazzara, Faisal Abouzaid, Nicola Dragoni, and Anirban Bhattacharyya. Toward design, modelling and analysis of dynamic workflow reconfigurations - a process algebra perspective. In *WS-FM*, pages 64–78, 2011.

[16] Manuel Mazzara and Anirban Bhattacharyya. On modelling and analysis of dynamic reconfiguration of dependable real-time systems. *DEPEND, International Conference on Dependability*, 2010.

[17] T. Mens, J. Magee, and B. Rumpe. Evolving software architecture descriptions of critical systems. *IEEE Computer*, 43(5):42–48, 2010.

[18] Angelo Morzenti and Pierluigi San Pietro. Object-oriented logical specification of time-critical systems. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 3(1):56–98, 1994.

[19] J.S. Ostroff. *Temporal Logic for Real Time Sytems.* Advanced Software Development Series. John Wiley & Sons, 1989.

[20] Matteo Pradella, Angelo Morzenti, and Pierluigi San Pietro. Refining real-time system specifications through bounded model- and satisfiability-checking. In *ASE*, pages 119–127, 2008.

[21] Matteo Pradella, Angelo Morzenti, and Pierluigi San Pietro. Bounded satisfiability checking of metric temporal logic specifications. *ACM Trans. on Soft. Eng. and Meth. (TOSEM)*, 2013.

[22] A. Robinson. *Non-standard analysis.* Princeton University Press, 1996.

[23] W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow patterns. *Distributed and Parallel Databases*, 14(1):5–51, July 2003.

[24] Moshe Y. Vardi and Pierre Wolper. An automata-theoretic approach to automatic program verification. In *Proc. of LICS*, pages 332–344, 1986.